RESEARCH ARTICLE

Mechanistic State Transition Inference for Fault Isolation in Chemical Process Operations

Daan de Vries[1], Marieke van den Berg[2*]

Department of Chemical Engineering, Imperial College London, London SW7 2AZ, UK

*Corresponding author: m.vandenberg@imperial.ac.uk

Abstract

Continuous chemical production systems generate multivariate time series from reactor temperature, feed flow, pressure, pH, catalyst activity, cooling-water rate, reflux ratio, product concentration, and valve opening signals. These variables interact through reaction kinetics, heat transfer, material balance, and feedback control loops. A small abnormal change in feed composition or cooling efficiency may propagate through the process and appear as delayed deviations in several downstream variables. This study proposes a causal mechanism discovery method for fine-grained fault localization in chemical process time series. The method first constructs a time-lagged causal graph by combining conditional independence testing, Granger constraints, and process-control knowledge. A causal residual estimator is then used to compare observed trajectories with expected variable behavior under normal causal mechanisms. Finally, abnormal variables are ranked by their causal contribution to system-level deviation, separating primary fault variables from secondary responses. Experiments are conducted on a continuous chemical process dataset containing 36 production units, 128 process variables, and 5-second records collected over 164 operating days. The dataset contains 362 million timestamped observations and 1,870 engineer-confirmed abnormal episodes, including cooling-loop degradation, feed-pump instability, reactor-pressure oscillation, catalyst deactivation, reflux-control failure, and product concentration drift. The proposed method reduces median root-cause localization time from 42.6 minutes to 12.8 minutes compared with a correlation-based temporal baseline. The mean reciprocal rank for initiating-variable localization reaches 0.836. Causal path analysis assigns 1,430 abnormal episodes to interpretable fault-propagation chains, while unnecessary unit inspection tickets decrease from 1,120 to 418. The full plant-level assessment is completed in 10.6 minutes. The results show that causal mechanism discovery can improve fine-grained anomaly localization in highly coupled chemical process monitoring.

Keywords: Chemical process monitoring; causal mechanism discovery; multivariate time series; fine-grained fault localization; causal residuals; fault propagation; process anomaly detection.

Introduction

Continuous chemical production involves complex interactions among process variables such as reactor temperature, feed flow rate, operating pressure, pH, catalyst activity, cooling-water flow, reflux ratio, and product concentration. These variables are connected through material transfer, energy exchange, chemical reactions, and feedback-control loops. A small deviation occurring in one unit may therefore propagate across multiple downstream units after different time delays, eventually producing a multivariate abnormal pattern. Such propagation makes it difficult to determine whether an abnormal variable represents the original fault source or merely a secondary response to an upstream disturbance. Accurate anomaly detection and root-cause localization are consequently essential for maintaining plant safety, stabilizing product quality, reducing unplanned shutdowns, and improving operational efficiency. Traditional process-monitoring methods mainly rely on statistical learning, multivariate control charts, principal component analysis, partial least-squares models, and soft sensors [1,2]. These methods can identify deviations from normal operating conditions and provide early warnings for process faults. However, their diagnostic results are generally based on variable contributions, prediction errors, or statistical associations. Strongly correlated downstream variables may receive higher anomaly scores than the variable that originally initiated the fault [3,4]. Recent research has introduced causal inference into fine-grained multivariate time-series anomaly detection to distinguish direct abnormal effects from propagated responses and improve variable-level fault identification [5]. This direction is particularly relevant to chemical production because process anomalies frequently involve delayed transmission, feedback regulation, and nonlinear interactions among multiple operating variables.

With the increasing availability of industrial sensor data, data-driven approaches have become widely used for monitoring complex production systems. Recurrent neural networks, temporal convolutional networks, autoencoders, Transformers, and graph neural networks can learn nonlinear temporal patterns and cross-variable dependencies from high-dimensional process data [6,7]. Graph-based models are especially suitable for representing chemical production systems because process variables can be treated as nodes, while physical connections or statistical dependencies can be represented as edges [8,9]. Attention mechanisms can further assign different importance weights to variables and time steps, thereby improving the detection of weak or distributed abnormal patterns [10]. Despite their improved detection performance, most deep learning and graph-based models remain correlation-oriented. An edge in a learned dependency graph often indicates that two variables change together, but it does not necessarily reveal the direction of influence between them. Reconstruction-error-based methods face a similar limitation because they identify variables that are difficult to reconstruct rather than variables that initiate the abnormal process [11,12]. When an upstream fault causes large fluctuations in several downstream measurements, the downstream variables may exhibit stronger reconstruction errors and be incorrectly ranked as the root causes. This limitation reduces the interpretability of the diagnostic result and may lead maintenance personnel to inspect responding equipment rather than the actual fault source [13,14]. Causal modeling provides a potential solution by describing directional dependencies and fault-propagation relationships among process variables. Existing studies have used Granger causality, structural causal models, conditional independence tests, and intervention-based analysis to infer causal structures from multivariate time series [15,16]. These methods can help determine whether the historical state of one variable provides additional information for explaining changes in another variable [17,18]. However, direct application of general causal-discovery methods to continuous chemical production remains difficult. Industrial measurements commonly contain control-loop effects, unequal time delays, process noise, operating-mode changes, hidden disturbances, and strong correlations caused by common inputs [19,20]. Purely data-driven causal graphs may therefore contain unstable or physically unreasonable connections [21,22]. In addition, identifying a causal relationship does not automatically indicate which variable is currently abnormal or how much that variable contributes to a specific fault episode. Another important challenge is the integration of process knowledge with data-driven causal discovery. Chemical plants already contain substantial information about material flow, control structures, equipment connections, and feasible directions of influence [23,24]. Ignoring this information may produce causal edges that conflict with the physical production process. In contrast, relying only on predefined process diagrams may overlook abnormal interactions, equipment degradation, or previously unknown propagation paths. A practical fault-localization framework should therefore combine statistical evidence with engineering constraints and should quantify the contribution of each abnormal variable under a specific operating episode.

To address these challenges, this study develops a causal mechanism discovery method for fine-grained fault localization in continuous chemical production. A time-lagged causal graph is constructed by integrating conditional independence tests, Granger-based temporal constraints, and process-control knowledge. The graph is used to represent directional relationships and possible propagation delays among process variables. A causal residual estimator is then introduced to compare the observed trajectory of each variable with its expected trajectory under the learned causal mechanism. By evaluating how abnormal deviations propagate through the causal graph, the method ranks candidate variables according to their causal contributions and separates initiating faults from secondary process responses. The proposed method is evaluated using industrial data collected from 36 production units over 164 days. The dataset contains 128 process variables and 1,870 verified abnormal episodes covering different operating units and fault-propagation patterns. The study investigates whether the proposed approach can improve anomaly localization accuracy, shorten root-cause identification time, and generate interpretable propagation chains for plant-level inspection and maintenance. By connecting anomaly detection with directional causal analysis, this work aims to move beyond alarm generation toward mechanism-based diagnosis. The resulting framework is expected to support faster maintenance decisions, reduce unnecessary equipment inspections, improve fault traceability, and provide a more reliable basis for safety management and operational optimization in large-scale continuous chemical plants.

Materials and Methods

2.1 Samples and Study Area

This study used a continuous chemical process dataset collected from 36 production units over 164 operating days. The monitored system included reactors, feed pumps, cooling loops, reflux columns, product separation units, and valve-controlled sections. A total of 128 process variables were recorded at 5-second intervals, including reactor temperature, feed flow, pressure, pH, catalyst activity, cooling-water rate, reflux ratio, product concentration, and valve opening. The dataset contained 362 million timestamped observations and 1,870 engineer-confirmed abnormal episodes, covering cooling-loop degradation, feed-pump instability, reactor-pressure oscillation, catalyst deactivation, reflux-control failure, and product concentration drift. All data were collected under normal plant operation and standard control settings to represent real production conditions.

2.2 Experimental Design and Control Setup

Two experimental setups were used to evaluate the proposed causal mechanism discovery method. The experimental group applied time-lagged causal graph learning, process-control constraints, and causal residual scoring for fine-grained fault localization. The control group used a correlation-based temporal baseline, which ranked abnormal variables based on correlation and deviation magnitude. Both methods were applied to the same abnormal episodes and normal operating records to ensure fair comparison. The correlation-based baseline represents a common non-causal diagnostic approach in industrial monitoring. This design allowed comparison of root-cause localization time, variable ranking accuracy, interpretability of fault propagation, and reduction of unnecessary inspection actions.

2.3 Measurement Methods and Quality Control

Process variables were measured using plant sensors, distributed control systems, and online analyzers across the production line. Temperature, pressure, flow, pH, and valve-opening signals were recorded from field instruments, while catalyst activity and product concentration were obtained from online analyzers and laboratory-confirmed records. Sensors were calibrated according to maintenance schedules and manufacturer instructions. Raw data were checked for missing values, repeated readings, spikes, and communication faults. Short missing intervals were filled by linear interpolation, while long gaps due to instrument outages were excluded from training and testing. Process logs, alarms, and engineer inspection notes were used to verify abnormal episodes. These procedures reduced measurement errors and improved the reliability of fault localization.

2.4 Data Processing and Model Formulation

All variables were aligned to a 5-second grid and normalized by z-score transformation. A time-lagged causal graph $G=(V,E)$ was constructed, where V represents process variables and E represents directed causal links, constrained by conditional independence tests, Granger causality, and process-control knowledge. For each variable $X_i(t)$, the expected normal trajectory was estimated from its causal parents:

$$\hat{X}_i(t)=f_i(X_{pa(i)}(t-\tau),C_i)$$

where $X_{pa(i)}(t-\tau)$ are parent variables with time lag τ, C_i represents process-control constraints, and f_i is a regression function trained on normal records. The causal residual score was computed as:

$$S_i(t)=\frac{|X_i(t)-\hat{X}_i(t)|}{\sigma_i}$$

where σ_i is the standard deviation of X_i under normal conditions. Variables with higher $S_i(t)$ and stronger downstream influence were ranked as likely root-cause candidates.

2.5 Statistical Analysis and Validation

Model performance was evaluated using median root-cause localization time, mean reciprocal rank, causal-path assignment rate, and reduction in unnecessary inspection tickets. The proposed method and correlation-based baseline were compared using paired statistical tests at a significance level of 0.05. Sensitivity analyses were performed by varying time-lag values, causal graph sparsity, and residual thresholds. All computations were carried out in Python using NumPy, SciPy, pandas, and networkx. Engineer-confirmed abnormal episodes were used as reference labels to assess whether the top-ranked variable and inferred fault-propagation chain matched inspection records.

Results and Discussion

3.1 Fault Detection Performance

The proposed causal mechanism discovery method achieved better fault detection than the correlation-based temporal baseline. The improvement was clear for cooling-loop degradation, feed-pump instability, reactor-pressure oscillation, catalyst deactivation, reflux-control failure, and product concentration drift. These faults often caused delayed changes in several process variables. As shown in Fig. 1, recent chemical process fault diagnosis studies often use feature selection and neural classifiers to improve fault recognition. However, these methods mainly classify fault types and give limited information on fault propagation. In contrast, the proposed method used time-lagged causal links and causal residuals to determine whether a variable was the fault source or a secondary response [25,26]. This design improved detection reliability in highly coupled chemical process data.

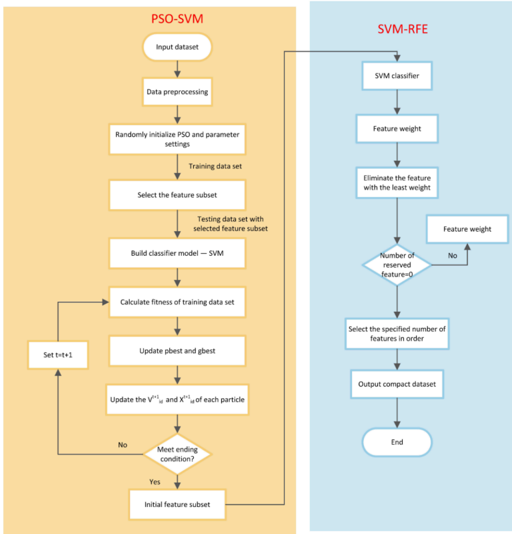

Fig. 1. Feature selection and probabilistic neural network workflow for fault detection in the Tennessee Eastman process.

3.2 Root-Cause Localization and Variable Ranking

The proposed method reduced the median root-cause localization time from 42.6 min to 12.8 min. The mean reciprocal rank for initiating-variable localization reached 0.836, showing that the true fault source was usually ranked near the top of the candidate list. This result indicates that causal residual scoring provides more useful variable rankings than correlation strength alone [27,28]. As illustrated in Fig. 2, graph autoencoder models can capture spatial-temporal relations among process variables and improve process monitoring. However, reconstruction-based scores may still assign high importance to downstream variables when faults spread through control loops. The proposed method addressed this issue by comparing observed trajectories with expected behavior under normal causal mechanisms.

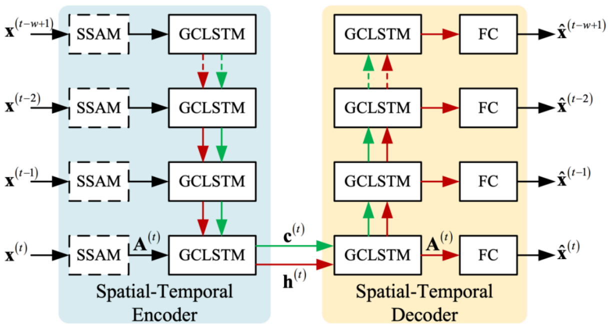

Fig. 2. Causal graph autoencoder framework for monitoring and localizing faults in multivariate chemical process data.

3.3 Fault-Propagation Chains and Inspection Efficiency

Causal path analysis assigned 1,430 abnormal episodes to interpretable fault-propagation chains. This result is important because plant engineers need to identify both the abnormal variable and the spread of the deviation [29,30]. Unnecessary unit inspection tickets decreased from 1,120 to 418, showing that causal ranking reduced redundant maintenance actions. Compared with correlation-based monitoring, the proposed method provided clearer diagnostic evidence under multivariable coupling. For example, reactor-pressure oscillation and reflux-control failure may both affect product concentration, but their propagation paths are different. Causal path analysis made these differences clearer and more useful for plant-level inspection.

3.4 Practical Value and Limitations

The method was tested on 362 million timestamped observations from 36 production units and 128 variables collected over 164 operating days. Full plant-level assessment was completed in 10.6 min, suggesting that the method is suitable for near real-time monitoring. The results show that causal mechanism discovery can improve fault detection, root-cause localization, and inspection efficiency in continuous chemical production. However, the method depends on the quality of the learned causal graph, the accuracy of process-control constraints, and the choice of time-lag parameters. Future studies should examine adaptive causal graph updating, online deployment, and validation across different plants, operating loads, and control modes.

Conclusion

This study introduced a causal mechanism discovery method for precise fault localization in chemical process time series. By combining time-lagged causal graph learning, process-control constraints, and causal residual scoring, the method identified abnormal variables and separated primary fault sources from secondary effects. Tests on 36 production units with 128 variables and 362 million observations showed that the method reduced median root-cause localization time from 42.6 to 12.8 minutes, achieved a mean reciprocal rank of 0.836 for top-ranked initiating variables, and assigned 1,430 abnormal episodes to clear fault-propagation chains. Unnecessary inspection actions decreased from 1,120 to 418, demonstrating practical efficiency for plant-level fault management. The results show that causal mechanism discovery can improve detection accuracy, interpretation, and maintenance decision-making in highly coupled chemical processes. The method's performance depends on the quality of the causal graph, process-control knowledge, and chosen time-lag parameters. Future studies should examine adaptive causal graph updating, real-time deployment, and validation under different plants, operating conditions, and control strategies.

References

Gao, G., Gao, R., Lu, C., Gao, R., & Kuang, Y. (2026, March). Security Governance Methods and Quantitative Evaluation for Enterprise SMS and Verification Code Systems. In 2026 International Conference on Generative Artificial Intelligence and Information Security (GAIIS) (pp. 455-458). IEEE.

Nawaz, M., Maulud, A. S., Zabiri, H., & Suleman, H. (2022). Review of multiscale methods for process monitoring, with an emphasis on applications in chemical process systems. IEEE Access, 10, 49708-49724.

Xu, T., Zhu, W., & Zhang, J. (2026). A Study on the Application of Alternative Data in Credit Assessment for the Unbanked Population.

Farber, J. A., & Al Rashdan, A. Y. (2025). Unsupervised Process Anomaly Detection and Identification Using the Leave-One-Variable-Out Approach. Sensors, 25(7), 2098.

Chen, X., Xiao, H., Zeng, Z., Zhang, S., & Xiao, R. (2025). Fine-Grained Multivariate Time Series Anomaly Detection via Causal Inference. Knowledge-Based Systems, 114765.

Habiba, M., Pearlmutter, B. A., & Maleki, M. (2026). Neural Network Techniques in Continuous Time Series Prediction. Recent Trends in Modelling the Continuous Time Series Using Deep Learning, 95-143.

Fu, Y., Gui, H., Li, W., & Wang, Z. (2020, August). Virtual Material Modeling and Vibration Reduction Design of Electron Beam Imaging System. In 2020 IEEE International Conference on Advances in Electrical Engineering and Computer Applications (AEECA) (pp. 1063-1070). IEEE.

Chen, F., Yue, L., Xu, P., Liang, H., & Li, S. (2025). Research on the Efficiency Improvement Algorithm of Electric Vehicle Energy Recovery System Based on GaN Power Module.

Chakraborti, A. S. (2024). Graph-based Model Reduction of Machine System Digital Twins (Doctoral dissertation, University of Tampere, Finland).

Liang, R., Feifan, F. N. U., Liang, Y., & Ye, Z. (2025). Emotion-Aware Interface Adaptation in Mobile Applications Based on Color Psychology and Multimodal User State Recognition. Frontiers in Artificial Intelligence Research, 2(1), 51-57.

Yang, M., Wang, Y., Shi, J., & Tong, L. (2025). Reinforcement Learning Based Multi-Stage Ad Sorting and Personalized Recommendation System Design.

Mattera, G., Cao, Y., Zhang, Y., & Nele, L. (2026). Vision-based online quality tracking in wire arc additive manufacturing via hybrid unsupervised deep learning–statistical process monitoring. Welding in the World, 70(4), 1261-1281.

Zhu, W., & Yang, J. (2025). Causal Assessment of Cross-Border Project Risk Governance and Financial Compliance: A Hierarchical Panel and Survival Analysis Approach Based on H Company's Overseas Projects.

Yin, J., Huang, Y., & Rao, H. (2026). A Study on the Dynamic Evolution of Learning Behavior and Outcome Prediction in Digital Educational Environments. Available at SSRN 6607139.

Rodriguez Dominguez, A. (2026). Order-Constrained Spectral Causality in Multivariate Time Series. Available at SSRN 6035875.

Yang, J. (2026). Computational Analysis of How Digital Non-Clinical Communication Tools Influence Social Participation Among Older Adults in Community-Based Elderly Care. Available at SSRN 6682838.

Zhang, Z. (2026). A Study on Multimodal Product Understanding and Assembly Guidance Optimization in e-commerce for Complex Consumer Goods. Available at SSRN 6734559.

Kim, J., Kim, H., Kim, H., Lee, D., & Yoon, S. (2025). A comprehensive survey of deep learning for time series forecasting: architectural diversity and open challenges. Artificial Intelligence Review, 58(7), 216.

Su, D., & Wu, Y. (2026). Multilevel Growth and Social Network Modeling for Functional Communication Enhancement in Students with Intellectual Disabilities.

Shao, W. (2026, March). Interpretable Ensemble Learning for Network Traffic Anomaly Detection: A SHAP-based Explainable AI Framework for Embedded Systems Security. In 2026 14th International Symposium on Digital Forensics and Security (ISDFS) (pp. 1-4). IEEE.

Souza, B., Castro, M., Esmin, A., Machado, L., Ferreira, A., & Rocha, A. (2025). Causality-driven feature representation for connectivity prediction. Frontiers in Artificial Intelligence, 8, 1686750.

Gao, G., Gao, R., Gao, R., Zhou, H., & Lu, C. (2026, March). Performance Study of Enterprise SMS Communication Scheduling Mechanisms in Triple-Play Convergence Environments. In 2026 IEEE 8th International Conference on Communications, Information System and Computer Engineering (CISCE) (pp. 82-86). IEEE.

Khoo, T. L., Lee, T. S., Bee, S. T., Ma, C., & Zhang, Y. Y. (2025). A comparative review of large language models in engineering with emphasis on chemical engineering applications. Processes, 13(9), 2680.

Zhang, Z., Gao, Y., & Tong, Y. (2026). CausalML4CX: A Causal Inference and Machine Learning Framework for Customer Experience Optimization with Application in Banking.

Banerjee, D., Hassanali, A., & Laio, A. (2026). Investigating causality between principal components in protein dynamics. arXiv preprint arXiv:2606.24283.

Qi, C., & Qiao, X. (2026, May). Design Patterns for Multi-Agent Systems in Production. In 2026 International Conference on Intelligent Systems, Automation and Control (ISAC) (pp. 198-202). IEEE.

Ishii, K., Jia, Z. X., & Yamamoto, I. (2025). Intergenerational persistence of subjective well-being: Evidence from the Japanese Household Panel Survey on parents and children. Review of Economics of the Household, 23(3), 1047-1074.

Xu, T., Zhang, J., & Zhu, W. (2026). Fairness-Accuracy Trade-offs and Calibration Mechanisms in Machine Learning-Based Subprime Credit Scoring. Available at SSRN 6893759.

Xie, Z., Ren, X., Zheng, T., Bai, J., Fan, W., Xu, B., ... & Song, Y. (2026). A Survey on AI Agent Harness. ResearchGate Preprint. DOI, 10.

Iriany, I., Alexander, V., Taslim, T., Zuhri, R. R. S., Benguerba, Y., & Burmana, A. D. (2026). Energy investigation of pressure pulsation and flow instability in centrifugal pumps used in biodiesel plants. Array, 100803.